

METHOD AND APPARATUS IN NETWORK MANAGEMENT SYSTEM FOR PERFORMANCE-BASED NETWORK PROTOCOL LAYER FIREWALL

Field of the Invention

This invention relates to distributed computing systems and more particularly to a system and method for managing the distribution of bandwidth at an endpoint of a distributed computing network.

Background of the Invention

Distributed data processing networks with thousands of nodes, or endpoints, are known in the prior art. The nodes can be geographically dispersed and the computing environment managed in a distributed manner with a plurality of computing locations running distributed kernels services (DKS). The managed environment can be logically separated into a series of loosely connected managed regions in which each region has its own management server for managing local resources. The management servers coordinate activities across the network and permit remote site management and operation. Local resources within one

region can be exported for the use of other regions in a variety of manners.

Managed regions within a highly distributed network may attempt to incorporate fault-tolerance with firewalls that attempt to limit any damage that might be caused by harmful entities. A firewall can prevent certain types of network traffic from reaching devices that reside on the "other" side, beyond the firewall. For example, a firewall can examine the frame types or other information of incoming data packets (i.e., so-called "packet sniffing") and decide to stop certain types of information that has previously been determined to be harmful, such as virus probes, pings, broadcast data, etc. Another use of such firewalls is to influence the distribution of bandwidth by denying access to certain types of communications which may unnecessarily consume needed bandwidth. Yet another role of a firewall is to prevent outside entities' attempts to breach an internal network (or network devices located beyond the firewall) to steal information and/or attack (i.e., "hack") the network. While existing firewalls can prevent certain entities from obtaining information from the protected network devices, firewalls can simultaneously present a barrier to the operation of legitimate, useful processes.

A firewall typically comprises a static dedicated piece of code that operates by using a dedicated port. Each software component communicates with another component by knowing the

dedicated port number of the other component. However, memory and other system constraints can eventually limit the number and the management of dedicated ports, and the dynamic reconfiguration of port numbers can be quite difficult. Another drawback to the static firewall system which is executed at the device driver level (i.e., the packet sniffing type firewall) is that the component must necessarily look at every packet which traverses that port. Given the quantity of communications in vast distributed networks, the analysis of every data packet can be an overwhelming task. If communications could be screened based on protocol, a significant amount of packet analysis could be foregone.

Yet another drawback to the presently available firewall technology is that it provides a "yes" or "no" approach to evaluating communications, whereby usage is either permitted or denied. There exists no mechanism today for a performance-based analysis of network communications at a firewall in order to allow continued usage provided that the bandwidth being consumed is within predetermined limits.

It is desirable, therefore, and is an object of the present invention, to have a method and apparatus for providing a performance-based firewall in a distributed network environment.

Another object of the present invention is to provide a firewall which can dynamically influence distribution of bandwidth in a network.

Yet another object of the present invention is to provide a firewall at the protocol layer rather than the packet layer.

Summary of the Invention

The foregoing and other objects are realized by the present invention wherein a method and apparatus are disclosed for implementing a performance-based firewall at the protocol layer. Application Action Objects (AAOs) are created for requesting applications and are mapped to specific protocol events. Each AAO is then used as a Usage Based Firewall (UBF) to monitor all usage of the protocol at the endpoint identified by the application, thereby acting as a performance-based, protocol layer firewall for communications at that endpoint. A responsible logical gateway monitors the AAO and reports AAO activity to a UBF Manager at a control server to direct the AAO regarding continued usage based on bandwidth considerations.

Brief Description of the Drawings

The invention will now be described in greater detail with specific reference to the appended drawings wherein:

Fig. 1 provides a general schematic diagram of a distributed network environment;

Fig. 2 provides a more specific schematic diagram of the components in a control server of the distributed network in accordance with the present invention;

Fig. 3 provides a representative process flow for implementing a protocol specific, usage based firewall at an endpoint in accordance with the present invention; and

Fig. 4 provides a representative process flow for evaluating the continued usage of a usage based firewall which has been deployed to control communications that flow to a network endpoint.

Description of the Preferred Embodiment

The present invention can be implemented in any network with multiple servers and a plurality of endpoints; and is particularly advantageous for vast networks having hundreds of thousands of endpoints in which an application would like to exercise endpoint control over communications. Fig. 1 provides a schematic illustration of a network for implementing the present invention. Among the plurality of servers, 101a-101n as illustrated, at least one of the servers, 101a in Fig. 1, which already has some of the distributed kernel services (DKS), is

designated as a Usage Based Firewall (UBF) control server for the purposes of implementing the present invention.

A network has many endpoints, with endpoint being defined, for example, as one Network Interface Card (NIC) with one MAC address, IP Address. The control server 101a in accordance with the present invention has the components illustrated in Fig. 2 for providing a method including the steps of: receiving an application request for a firewall at a network endpoint; creating an Application Action Object in response to a request from an application which wishes to control what flows to the endpoint; registering the protocol request and obtaining a session number for the AAO from the UBF Manager at the control server; returning or deploying the AAO to the requesting application to act as the UBF for the endpoint; and, monitoring activities at the AAO and responding thereto.

Fig. 2 illustrates in greater detail the control server components which are relevant to the implementation of the present invention. Additional core server components and their functionality, as have been detailed in co-pending application entitled "METHOD AND SYSTEM FOR MANAGEMENT OF RESOURCE LEASES IN AN APPLICATION FRAMEWORK SYSTEM", Serial No. _____, filed _____, the teachings of which are incorporated by reference herein (Docket AUS9-2000-0699), are not repeated in detail in this description, since those components and their functionality do not change for the specific implementation of the present

invention being described herein. As shown in Fig. 2, the server 200 includes the already-available DKS core services at component 202, which services include the object request broker (ORB) 212, service manager 222, and the Administrator Configuration Database 232, among other standard DKS services. The ORB 212 will create the Application Action Objects (AAOs) in response to application requests to the server and pass those AAOs to sit at the specified endpoints, as further detailed below. The Administrator Configuration Database 232 will include stored definitions for the allowable protocol for endpoints and the endpoint addresses, along with endpoint-specific usage values (e.g., maximum numbers for requests for an endpoint per minute and/or per hour) or default values, for use in performance-based analysis (discussed below) when administrating the firewall in operation.

The DKS Internet Protocol Object Persistence (IPOP) Manager 203 provides the functionality for gathering network data, as is detailed in the aforementioned co-pending patent application, along with an Application Action Object (AAO) Decoder for endpoints 223, discussed below, and a Protocol Usage Based Firewall (UBF) Database 213 for endpoints, the latter two components being specific to the present invention. The inventive role of the IPOP Manager components is to keep track of metrics to estimate bandwidth at an AAO deployed at an endpoint and to compare those monitored values to system

administrator-defined thresholds (stored at the Administrator Configuration Database 232). In addition to the enhanced IPOF Manager 203, the server of the present invention includes a Usage Based Firewall (UBF) Manager 204, the functions of which are further detailed below. The UBF Manager 204 includes a UBF Mapper 214 for mapping Application Action Objects (AAOs) to protocol events and a Database 224 comprising Protocol Session Counter per Application data for use as further discussed below.

The specific roles of the relevant components of Fig. 2 will become apparent in the following discussion of the operations of the present invention. Because distributed kernel services are available in the distributed network, the present system can control at which endpoint a so-called "traveling" firewall is placed. Furthermore, the present invention allows the traveling endpoint firewall to be protocol-specific such that, rather than implementing a generic "GetSocket(...)" command as in the prior art to statically deploy a packet layer firewall, the present invention can effectively implement a "GetFTPsocket", "GetPingSocket", or other protocol-specific command at an endpoint to act as a protocol-specific firewall. Finally, the inventive firewall is an "interactive" session object, the actions or performance of which can be monitored to prevent overuse of the endpoint.

Fig. 3 provides a representative process flow for implementing the protocol-specific, interactive, usage-based

firewall at an endpoint in accordance with the present invention. When an application wishes to control what flows to a specific endpoint, the application will request an Application Action Object (AAO) from the ORB at the server in step 301. The request is handled by a logical DKS Gateway (not shown) which asks the IPOP Manager to decode the endpoint at step 303. In order to decode the endpoint, the DKS Gateway takes a Object Identifier (IPOPOid) and determines the physical network address of the target endpoint in addition to determining which DKS Gateway(s) will be used to route this action object request. The decoded information from the AAO Decoder 223 of IPOP Manager 203 is provided to the ORB to be added to the AAO at step 305. Next, for requests in which the protocol is specified, the IPOP Manager registers the protocol request with the UBF Manager 204 at step 307. If the request does not specify a protocol, the UBF Mapper is used to determine the protocol for the request (e.g., a "move" application action object request would invoke the use of the File Transfer Protocol (FTP)) prior to registering the protocol.

At step 309, the UBF Manager adds a session number to the AAO, which session number will be used for monitoring all usage of the protocol by the UBF Manager. Thereafter, the AAO with session number is returned to the IPOP Manager at step 311. The IPOP returns the AAO to the logical DKS Gateway at 313, followed by returning the AAO to the application at step 315.

In operation, the AAO will be used by the application and its use will optimally be monitored for performance-based analysis. Fig. 4 provides a representative process flow for evaluating the continued usage of a usage based firewall which has been deployed to control that which flows to a network endpoint. When an application uses the AAO, for example to ping the endpoint, at step 401, the use comprises executing an action method which initiates routing of the AAO to a responsible gateway at 403. The gateway, in turn, notifies the UBF Manager at step 405 that the protocol has been used. The gateway uses the session number which is in the AAO when notifying the UBF Manager so that the AAO is appropriately identified. Again using the session number, the gateway asks the UBF manager at 407 if continued usage of the AAO at the endpoint is permissible. The UBF Protocol Session Counter Database 224 will retrieve the configured maximums for requests for the protocol/application combination, obtain the current count of requests for the protocol/application combination from the Protocol Session Counter 224, and will compare the current count to the configured maximum. If the current count does not exceed the configured maximum, then the UBF will notify the gateway that it may perform the action at the endpoint at 408. If the configured maximum is exceeded, such that the determination at decision box 407 is that continued use is impermissible, the application will be informed of the overuse at step 409. In addition to notifying the

